

Docket No. 157923-0019  
Express Mail Label No. EL593465073US

UNITED STATES PATENT APPLICATION

FOR

**METHOD AND APPARATUS FOR MANAGING PAYLOAD  
BUFFER SEGMENTS IN A NETWORKING DEVICE**

INVENTORS:

Salil Suri  
David Geddes  
Scott Furey  
Michael Moretti  
Thomas Wu

PREPARED BY:

IRELL & MANELLA LLP  
840 Newport Center Drive  
Suite 400  
Newport Beach, California 92660  
(949) 760-0991

## CROSS REFERENCE TO RELATED APPLICATIONS

This application is related to and claims priority from provisional application serial number 60/441,851, entitled "Method and Apparatus for Managing Payload Buffer Segments in a Networking Device," filed on January 21, 2003.

5

## BACKGROUND OF THE INVENTION

### 1. Field of the Invention

The present invention relates in general to data networks and more particularly, to a method and apparatus for managing payload buffer segments in a networking device.

### 10 2. Background of the invention

Fibre Channel is a computer communications protocol designed to provide for higher performance information transfers. Fibre Channel allows various existing networking protocols to run over the same physical interface and media. In general, Fibre Channel attempts to combine the benefits of both channel and network technologies.

A channel is a closed, direct, structured, and predictable mechanism for transmitting data between relatively few entities. Channels are commonly used to connect peripheral devices such as a disk drive, printer, tape drive, etc. to a workstation. Common channel protocols are Small Computer System Interface (SCSI) and High Performance Parallel Interface (HIPPI).

20 Networks, however, are unstructured and unpredictable. Networks are able to automatically adjust to changing environments and can support a larger number of connected nodes. These factors require that much more decision making take

place in order to successfully route data from one point to another. Much of this decision making is done in software, making networks inherently slower than channels.

- 5      Fibre Channel has made a dramatic impact in the storage arena by using SCSI as an upper layer protocol. Compared with traditional SCSI, the benefits of mapping the SCSI command set onto Fibre Channel include faster speed, connection of more devices together and larger distance allowed between devices. In addition to using SCSI, several companies are selling Fibre Channel devices that run Internet Protocol (IP).
- 10     Fibre Channel will most likely continue to expand into the storage markets, which will make use of its benefits over traditional channel technologies such as SCSI. Being able to access mass storage devices quicker and from greater distances is very attractive to such applications as multimedia, medical imaging, and scientific visualization. As such, there is a need for an improved system for processing information in the Fibre Channel arena.
- 15

## BRIEF SUMMARY OF THE INVENTION

Methods and apparatus for managing payload buffer segments in a networking device are disclosed. One method comprises writing a plurality of data frames that were received over a first network interface according to a source protocol to select locations of a plurality of data segment locations of a payload buffer, wherein each of the select locations is based on a current index value for a current position of a read pointer of a FIFO buffer. In this method the FIFO buffer has the read pointer, a write pointer, and a plurality of entries that stores a plurality of index values corresponding to the plurality of data segment locations. The method further comprises incrementing the read pointer to a subsequent entry of the plurality of entries after writing each of the plurality of data frames, and transmitting, in accordance with a destination protocol, the plurality of data frames over a second network interface in a different order than the plurality of data frames were received over the first network interface.

15        Other embodiments are disclosed and claimed herein.

## BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1A-1B illustrates a block diagram of one embodiment of an ASIC capable of carrying out one or more aspects of the present invention.

Figures 2A-2E illustrate one embodiment of the interaction between a Free  
5 Pool FIFO and a corresponding payload buffer.

Figure 3 depicts one embodiment of the ingress and egress payload buffers of the ASIC of Figures 1A-1B.

Figure 4 illustrates one embodiment of a process for managing payload buffer segments consistent with the principles of the invention.

## DETAILED DESCRIPTION OF EXEMPLARY EMBODIMENTS

One aspect of the invention relates to processing frames in an order that may be different than the order in which they were received. In one embodiment, a Free Pool FIFO is used to deliver higher priority frames ahead of the lower priority frames. In another embodiment, delivering higher priority frames ahead of lower priority frames is done using a single data queue that is associated with the Free Pool FIFO.

5

Another aspect of the invention is to enable a networking device to have the capability of reordering out-of-order frames. In one embodiment, this is done by enabling the networking device to hold an out-of-order frame in a buffer until the expected in-order frame arrives. Again, a Free Pool FIFO as disclosed herein may be used to cause release the frames in a different order than the order in which they were received. In one embodiment, the reordering of out-of-order frames is accomplished using an on-chip processor.

10

15 It should be appreciated that the frame processing operations performed by the Free Pool FIFO and corresponding payload buffer may be in a protocol bridge that is receiving FC frames which are to be sent out over a POS interface. However, it should similarly be appreciated that the principles of the invention may similarly be applied in other applications.

### **20 I. SYSTEM OVERVIEW**

#### **A. Hardware Design**

Figures 1A-1B illustrate a block diagram of one embodiment of an ASIC 10 capable of carrying out one or more aspects of the present invention. In the embodiment of Figures 1A-1B, the ASIC 10 includes two Fibre Channel (FC) ports,

F0 Port and F1 Port, with hardware associated with the F0 Port residing on the F0 function level and hardware associated with the F1 Port residing on the F1 function level. It should be appreciated, however, that there may be more or fewer FC ports and one or more of the hardware components for different FC functions may be

5 integrated onto the same function level.

Ingress (Ingrs) and egress (Egrs) references in Figures 1A-1B describe the data path direction between the Packet-over-SONET Physical Layer (POS/PHY – abbreviated as POS hereafter) interface 12 and the Fibre Channel 14. However, while Figures 1A-1B and the following description are directed to sending and

10 receiving data between a Fibre Channel interface and a POS interface, it should equally be appreciated that the principles of the invention may similarly be applied to other network protocols and other applications. For example, rather than having a POS interface 12 coupled to a POS network, the interface may be a System Parallel Interface (a/k/a System Packet Interface), Utopia or the interface marked by AMCC

15 Inc. under the name FlexBUS™. Similarly, rather than having a Fibre Channel interface coupled to Fibre Channel 12, ASIC 10 may be interfaced to an IEEE-1394, Infiniband, and/or iSCSI network. However, for brevity the following discussion with refer to only POS networks and Fibre Channel.

The Network Processor 16 may be any processor to which the ASIC 10

20 interfaces through the POS interface. The Egress POS Internal Queue (EPIQ) 18 may contain headers of frames received from the POS interface 12. In one embodiment, POS frames that will be processed by the internal embedded processor (PRC) 20 are routed to the EPIQ 18. Similarly, the Egress POS Pass Through Queue (EPPQ) 22 may contain headers of POS frames received from the POS interface, where the

25 payloads for such POS frames are intended to pass through the ASIC 10 to Fibre

Channel 14. In the embodiment of Figure 1B, both EPIQ 18 and EPPQ 22 are components of Header Queue Memory (HQM) 24.

Continuing to refer to Figures 1A-1B, the Ingress POS Internal Queue (IPIQ) 26 may contain headers of POS frames that have been generated by PRC 20. In addition, the Ingress POS Pass Through Queue (IPPQ) 28 may contain headers for POS frames whose payloads were received from the Fibre Channel 14. Ingress Fibre Internal Queue (IFIQ) 30, as shown in Figure 1B, may contain headers of frames received from the Fibre Channel 14. In one embodiment, FC frames whose payloads will be processed by the PRC 20 may be routed to the IFIQ 30. Moreover, Ingress Fibre Pass Through Queue (IFPQ) contains headers of frames received from the Fibre Channel 14, according to one embodiment. FC frames whose payloads will pass through the ASIC 10 to the POS interface 12 may be also be routed to the IFPQ 30.

In the embodiment of Figure 1B, the Egress Fibre Internal Queue (EFIQ) 34 may contain headers of FC frames that have been generated by the PRC 20. In that case, the frames may be sent out on the Fibre Channel 14. Moreover, the Egress Fibre Pass Through Queue (EFPQ) 36 contains headers of FC frames whose payloads were received from the POS interface 12, according to another embodiment.

In one embodiment, the memory queues of HQM 24 (e.g., EFPQ 36, EFIQ 34, EPPQ 22, EPIQ 18, IFIQ 30, IFPQ 32, IPIQ 26, and IPPQ 28) are shared dual-port RAM that is accessible by the ASIC 10 hardware logic as well as PRC 20.

The Egress POS Control (EPC) 48 module may be used to provide read functionality to transfer data from the Network Processor 16 (or associated memory) to the Egress Payload Buffer (EPB) 40 module or to the Egress POS queue memory of HQM 24. Similarly, the Ingress POS Control (IPC) 50 module may be used to

provide the DMA write function to transfer data to the Network Processor 14 (or associated memory) from the Ingress Payload Buffer (IPB) 38 module or the Ingress POS queue memory of HQM 24.

The IPB 38 of Figure 1B may contain payloads for frames that will be sent to 5 the POS Interface 12. It should be appreciated that the payloads may have come from the Fibre Channel 14 or may have been created internally by the PRC 20. Moreover, the EPB 40 may contain payloads for frames that will be sent out on the Fibre Channel 14, where the payloads may either have come from the POS interface 10 12, or may have been created by the PRC 20. In one embodiment, the IPB 38 and EPB 40 modules are each comprised of a payload buffer and FIFO that may be used to carry out one or more aspect of the present invention. In particular, the IPB 38 and EPB 40 modules may be used, as will be described in more detail below in 15 Section II, to provide out-of-sequence release of payload buffer segments. In another embodiment, delivering higher priority frames ahead of lower priority frames may be done using one or both of these data queues. Similarly, the IPB 38 and EPB 40 modules may also provide the capability of reordering out-of-order frames by, for example, holding an out-of-order frame until the expected in-order 20 frame arrives.

The Fibre Channel interface provides the interface and control between the 20 Fibre Channel and the ASIC 10. In the embodiment of Figures 1A-1B, the Fibre Channel interface consists of 4 major modules – the Egress Fibre Channel Control (EFC) 44, Arbitrated Loop Control (ALC) 45, Ingress Fibre Channel Control (IFC) 46 and Fibre Channel Interface (FCI) 52 modules. In particular, the EFC module 44 may be used to provide the frame flow control mechanism of the FC transmitting 25 port (i.e., F0 or F1), while other operations which may be performed by the EFC module 44 include frame assembly, CRC generation, and retransmission of certain

data from the ALC module 45 (e.g., L\_Port data). In one embodiment, the EFC module 44 assembles and transmits frames to the FCI module 52 based on the data from HQM 24, EPB 40, and the ALC module 45.

In the embodiment of Figure 1A, the ALC module 45 is located between the IFC module 46 and EFC module 44. In one embodiment, this module consists primarily of a Loop Port State Machine (LPSM) whose main function is to continuously monitor the data stream coming from the IFC module 46. The LPSM may further be used to monitor commands from the PRC 20 and the EFC module 44. In one embodiment, the EFC 44 may send a command to the LPSM which defines the function to be performed by the ALC module 45 such as loop arbitration, open loop, close loop, etc. In another embodiment, the LPSM may be controlled by the PRC 20.

In one embodiment, the ALC module 45 may be used to detect different primitive signals or sequences (e.g., LIP, LPE, LPB, MRK, NOS, OLS, LR and LRR) and respond accordingly. In the loop topology, data from the IFC module 52 may be either passed on to the EFC module 44, or substituted with a primitive sequence depending on the function to be performed. The substitution may be either by the state machine itself or signaled from the EFC module 44.

The IFC module 36 may receive a data stream from the FCI module 52 and provides functions that may include frame disassembling, frame header matching and routing, FC\_FS primitive signal and sequence detection, CRC checking and link interface integrity measurement. In one embodiment, the data received from the FCI module 52 is passed on to the ALC module 45 for retransmission during a private/public loop (L\_Port) monitoring state. When not in the monitoring state, each frame received may be examined and routed to the appropriate destination

modules. If the frame has a payload, the payload may be written into the next available buffer segment in the IPB module 38, according to one embodiment.

The Processor Bridge Controller (PBC) module 54 provides the interfaces that connects the embedded processor (e.g., PRC 20) to the rest of the ASIC 10 hardware.

- 5 In the embodiment of Figure 1B, PRC 20 is coupled to the PBC module 54 via a PIF bus, which may be a general purpose I/O bus that supports burst reads and writes, as well as pipelined single-access reads and writes. In another embodiment, PRC 20 can also use the PBC module 54 to interface with external memory devices such as DDR/SDRAM 56 and NVRAM 58 attached to the ASIC 10 through the Memory Port I/F (MPI) module 60, or EEPROM 62 through the Initialization and Configuration Control (ICC) module 64. In yet another embodiment, the PBC module 54 may also provide bi-directional bridging between the F\_LIO bus 42 and Host Local I/O (H\_LIO) bus 66. In one embodiment, F\_LIO bus 42 may be used to provide access to registers in other hardware blocks through arbitration.
- 10
- 15 As previously mentioned, the MPI module 60 may be used to provide arbitrated accesses to external memory (e.g., DDR SDRAM 56 and/or NVRAM 58) devices by the PRC 20, as well as to every bus master on the internal H\_LIO bus 66.

- 20 In one embodiment, the ICC module 64 includes a Serial Memory Control (SMC) module, which can be used to initialize internal registers and provide read/write access to EEPROM 62. The ICC 48 may also include a trace control module (not shown) to provide external visibility of the internal signals.

## B. Frame Egress

In the embodiment of Figure 1A-1B, each frame that is received from the POS interface 12 may be routed to one of the two FC function levels (F0 or F1). As

mentioned previously, there may be more or fewer than two FC function levels, in which case the frames received from the POS interface 12 would be routed to whatever number of available FC function levels there may be. In one embodiment, frames are routed based (at least in part) on a port routing byte in a given frame header. In one embodiment, the port routing byte is located in the third byte of the frame header, although it should of course be understood that the port routing byte may be located elsewhere in the frame.

After the frame arrives at the selected function (e.g., F0 or F1 in this embodiment), a second routing decision may then be made based on a path routing bit. In one embodiment, the path routing bit is located in the POS frame header, and may be located in one of the first four bytes of the POS frame header. The path routing bit may be used to determine whether the frame will be routed to the "Pass-Through Path" or to the "Internal Path," where the Pass-Through Path is for frames containing payloads that are going to be sent out on Fibre, and the Internal Path is for frames whose payload contains configuration or control information that will be used by the PRC 20 and not sent out on Fibre.

In one embodiment, after the above-described routing decisions have been made, the received frame header is stripped from the payload and is stored in an entry in an Egress POS Queue (such as EPPQ 22 or EPIQ 18) that is dedicated to the selected function/path. A programmable number of bytes from the payload may also be stored along with the header. The payload may then separated from the frame and stored in the next available segment of the EPB 40 for the given FC function (F0 or F1). As will be described in more detail below in Section II, the EPB 40 module may be used to provide out-of-sequence release of payload buffer segments and/or delivery of higher priority frames ahead of lower priority frames.

A handle indicating which payload segment was used is stored by hardware in the HQM 24 queue which received the POS frame header.

In the case where the frame was routed to the Pass-Through Path, a portion of the frame header may be compared with the corresponding bytes from the previous 5 frame's header. If the contents of the bytes are equal, a 'header match' bit in the HQM 24 entry may be set indicating that the frames belong to the same context. It should be noted that the location of the bytes to be compared may be programmable via a bit mask. At this point, the PRC 20 may be notified that a frame has been received, while in another embodiment the PRC 20 is notified before the entire 10 payload has been received.

It should be appreciated that the PRC 20 may undertake a variety of operations at this point which may dependent upon several factors, including the path and contents of the frame, whether initialization has been completed, and in the case of an FCP frame, whether a command context already exists. Moreover, the 15 PRC 20 may undertake a frame Pass-Through operation and/or an Internal Frame operation, as will now be described.

### 1. Pass-through Frame Operation

As mentioned previously, a given frame may be routed to a Pass-Through Path or an Internal Path, depending on its path routing bit. Where the frame was 20 routed to the Pass-Through Path, the PRC 20 may be used to write the information necessary to create a suitable FC frame header. In one embodiment, the FC frame header is created in the next available entry in the EFPQ 36, although it may also be stored elsewhere. In one embodiment, the PRC 20 may also copy the payload segment handle to this EFPQ 36 entry. Moreover, if the frame belongs to the same 25 context as the previous frame, a bit may be set in the HQM 24 entry (e.g., EFPQ 36

entry) that instructs the hardware to automatically generate portions of the FC header based on values from the most recent FC frame that was generated from that queue.

After the PRC 20 has finished setting up the outgoing frame header, control  
5 of the HQM 24 entry may then be turned over to the hardware by setting a bit in the entry's control word. In one embodiment, this bit is used to determine whether the hardware or the software 'owns' the entry (e.g., has control of the entry). Similarly, other methods for releasing the entry may also be used. Once control of the HQM  
10 24 entry has been turned over to the hardware, the entry may then be queued up for transmission from one of the FC Ports. In one embodiment, frames that are released to the hardware are sent out on the FC Ports in the order in which they were released by the PRC 20. However, it should be appreciated that frames may be sent out in any number of other orders.

After the PRC 20 has set up an outgoing entry in the EFPQ 36, it may release  
15 the entry in the incoming EPPQ22 . In one embodiment, the entry is released by resetting a bit in the control word of the entry. Once released, the entry location may be reused for another egress POS frame header.

When the entry in the EFPQ 36 reaches the head of an HQM 24 queue, the hardware may automatically assemble an FC frame and send it out on the Fibre  
20 Channel 14, according to one embodiment. According to another embodiment, when this has been completed the hardware puts the completion status of the operation into the EFPQ 36 entry, and turns the entry over to the software. The EPB 40 segment may be returned to the free pool, or it may be returned by the PRC 20 after it checks the completion status in the HQM 24 entry.

## 2. Internal Frame Operation

If, on the other hand, the frame was routed to the Internal Path, the payload may be intended for use by the PRC 20. A programmable number of payload bytes may be made available to the PRC 20 in the entry in the EPIQ 18. In one embodiment, the EPIQ 18 may be made available to the PRC 20 in zero-wait-state memory. Moreover, additional payload bytes may be made available to the processor via the F\_LIO bus 42 (e.g., F0\_LIO and F1\_LIO).

After the PRC 20 has finished processing the information from the frame, it may release the entry in the EPIQ 18 to the hardware by resetting a bit in the control word of the entry. In one embodiment, the PRC 20 returns the payload buffer segment to the free pool by writing a segment handle to the payload segment release register.

## 3. Special Payload Buffer

If the PRC 20 needs to generate an egress FC frame, in one embodiment it may do so using the EFIQ 34 and a Special Payload Buffer (not shown). In one embodiment, the Special Payload Buffer is a single segment buffer consisting of 512 bytes and resides in zero-wait-state processor memory. After the PRC 20 has put the required information into the HQM 24 entry (e.g., in the EFIQ 34 entry) and Special Payload Buffer, the frame may then be released to the hardware by setting a bit in the HQM 24 entry, causing the frame to be sent out when the entry reaches the head of the particular queue.

## 4. Optional Headers

When a POS frame is received, its payload may be placed into an entry in the EPB 40. For Pass-Through payloads, the PRC 20 may occasionally be required to

insert an optional FC header between the FC header and the payload received from the POS interface 12. In order to accommodate this, a predetermined number of bytes may be allocated in each entry in the egress FC Header queues (e.g., EFPQ 36 and EPPQ 22). In one embodiment, the predetermined number of bytes is 72 bytes.

5 When the PRC 20 needs to insert an optional header, it writes the header to one or more of these spare byte locations in the HQM 24 entry, according to one embodiment. In addition, the PRC 20 may write the length of the optional header to a field (e.g., imm\_datafld\_size field) of the HQM 24 entry. Once the given HQM 24 entry has been turned over to the hardware and has reached the head of the queue,

10 the entry may be sent out to the Fibre 14. In one embodiment, the FC header is sent out first, followed by the bytes containing the optional FC header, followed by the payload. If multiple FC frames are generated from one entry in an FC Header queue, the hardware may be configured to include the optional header in each FC frame, or alternatively, in only the first frame.

15        5.     Raw Frames

Raw FC frames may be received from the POS interface 12 and sent out on the Fibre Channel 14 using the same process used with Pass-through frames described above in Section I.B.1. POS frames containing encapsulated raw FC frames may be routed to the Pass-Through path. In one embodiment, the POS frame header is stripped off and is placed into an entry in the EPPQ 22, while the encapsulated FC raw frame is automatically placed into the next available segment 20 of the EPB 40.

After the PRC 20 has been notified of the arrival of the POS frame, it may then perform the steps described above in Section I.B.1., except that a bit may be set in the 25 EFPQ 36 that direct the system to take most the information needed to build the FC

frame header from the raw FC frame in the EPB 40, rather than from the HQM 24 entry. In one embodiment, when this bit is set, the only fields that are taken from the HQM 24 entry are the SOF and EOF characters, and the S\_ID and D\_ID (i.e., Source-ID and Destination-ID, respectively). The remaining FC header fields may  
5 then be taken directly from predefined locations in the raw frame in the EPB 40.

Additional bits in the HQM 24 entry may be used by the PRC 20 to determine which mechanism will be used to generate the CRC (“Cyclic Redundancy Check”) checksum for the Fibre Channel 14 frame. In one embodiment, the possible mechanisms include: a) using the checksum located in the raw frame in the EPB 40,  
10 b) using a hardware generated checksum in the place of the one located in the EPB 40, and c) appending a hardware generated checksum to the end of the data in the EPB 40.

## 6. Cut-Through and Store-Forward Modes

In one embodiment, ASIC 10 provides two modes of operation. With the first  
15 mode, referred to herein as the Store-Forward mode, frames are received in their entirety from the POS interface 12 before they are sent out on the Fibre Channel 14. Alternatively, a Cut-Through Mode may be used, as described in co-pending United States Patent Application serial number 10/435,214, entitled “Method and Apparatus for Implementing a Data Frame Processing Model,” filed on May 8, 2003,  
20 the contents of which are hereby incorporated by reference. After a frame header and a programmable number of payload bytes have been received from the POS interface 12 in this mode, the frame may be output on the Fibre Channel 14. Thus, receiving and sending operations may overlap. In one embodiment, Cut-through mode may be enabled on a frame-by-frame basis.

## 7. Small FC Frames

Some Fibre Channel devices may negotiate a maximum FC payload size that is less than a nominal size, which in one embodiment is just over 2KB. In one embodiment, this negotiated size may be 512 bytes, although other sizes may also be 5 negotiated. In such a case, ASIC 10 may allow the Network Processor 16 to send nominal sized POS frames (e.g., 2KB) to the ASIC 10 for such devices, but will segment the POS frame into multiple FC frames to accommodate the smaller negotiated FC payload size.

When a POS frame is received by the ASIC 10, the header and payload may 10 be separated and routed to the EPPQ 22 and EPB 40 in the same manner described above for Pass-Through operations. In order to accommodate the smaller negotiated FC payload size, when the PRC 20 sets up an outgoing FC frame header in the EFPQ 36, it may indicate the negotiated size of the FC payload for a given device in the field in the HQM 24 entry (e.g., the 'maximum-send-size' field).

15 By way of a non-limiting example, the maximum-send-size field may be programmed with a value of 512 bytes instead of the nominal value of 2K. The remainder of the fields in the FC HQM 24 entry may then be filled in by the PRC 20 in the usual manner, after which the entry is released to the hardware. When the entry in question in the EFPQ 36 reaches the head of the queue, the value in the 20 'maximum-send-size' field may be compared to the value in another field (e.g., the 'expected-payload-size' field) of the same entry. If the 'expected-payload-size' field is larger, the system will generate multiple Fibre Channel frames. While in one embodiment, the generated multiple FC frames each have the payload size indicated by the 'maximum-send-size' field, it should be appreciated that they may also have 25 smaller payload sizes. In one embodiment, the generated FC frame use information

from the original HQM 24 entry, while in another embodiment, the hardware automatically increments certain fields in the subsequent FC headers, such as the SEQ\_CNT and Relative Offset fields.

Moreover, if the FC HQM 24 entry indicates that the data contained in the  
5 payload is the last data in an FC sequence, or that the FC Sequence Initiative should be transferred, the appropriate bits may be set in the header of only the last FC frame that is generated.

#### 8. Jumbo Frames

Another aspect of the invention is for the ASIC 10 to be configurable to accept  
10 normal frames, jumbo frames, or an intermix of normal and jumbo frames from the POS interface 12. For purposes of the present discussion, a normal frame is defined as a frame whose payload can fit into a single segment of the EPB 40, while a jumbo frame is a frame whose payload spans two or more segments of the EPB 40. In one embodiment, the maximum size of a jumbo frame is configurable up to a maximum  
15 of 32K bytes.

When a jumbo frame is received on the POS interface 12, the system may automatically allocate the necessary number of EPB 40 segments to hold the frame. Also, the system may allocate an entry in the EPPQ 22 for each EPB 40 segment that is allocated. These additional HQM 24 entries do not contain copies of the POS  
20 header, according to one embodiment. Instead, they may merely contain a pointer to a EPB 40 segment and indicate that the buffer segment contains overflow data belonging to the previous entry(ies) in the POS queue of the HQM 24.

While a jumbo frame is being received on the POS interface 12, the POS HQM 24 entries that are associated with each new EPB 40 segment may be turned over to

the processor incrementally as each EPB 40 segment is allocated. In one embodiment, each time the PRC 20 receives a POS HQM 24 entry, it sets up an entry in the FC queue of the HQM 24, copies the EPB 40 segment handle to it, and turns the FC HQM 24 entry over to the hardware. Using this mechanism, the hardware

5 may send an FC frame containing the first portion of a jumbo frame payload out on the Fibre 14 while the remainder of the jumbo frame payload is still being received on the POS interface 12. Similarly, the EPB 40 module may be comprised of a payload buffer and FIFO that may be used to provide out-of-sequence release of payload buffer segments and/or delivery of higher priority frames ahead of lower

10 priority frames.

Since all of the FC frames generated from a jumbo frame will typically belong to the same context, the system is only required to set up a full FC header for the first FC frame. In one embodiment, the hardware may be programmed to automatically generate the FC headers for each subsequent FC frame based on

15 information from the preceding frame, as described in co-pending United States Patent Application serial number 10/434,872, entitled "Method and Apparatus for Controlling Information Flow Through a Protocol Bridge," filed on May 8, 2003, the contents of which are incorporated herein by reference.

If the final FC frame generated from a jumbo frame will be required to

20 transfer the FC Sequence Initiative, or to end a sequence, the PRC 20 should know in advance what the overall length of the jumbo frame will be. In one embodiment, this may be accomplished by including a frame size field in the header of the POS jumbo frame.

## 9. Arbitration

In one embodiment, egress FC Frames may originate in either the EFPQ 36 or the EFIQ 34. At any point in time, there may be multiple FC frame headers in each of these queues waiting to go out on the wire.

5        Within each queue, FC frames will be output in the order in which they were released to the hardware, according to one embodiment. However, the same principle need not apply between queues. For example, frames that are waiting in one queue may be delayed while newer frames in the other queue go out on the Fibre 14.

10      In one embodiment, the arbitration algorithm has two settings: ‘ping-pong’ and ‘sequence’. When the arbiter is programmed for ping-pong mode, egress FC frames may be taken from the EFPQ 36 and the EFIQ 34 in alternating order, one at a time from each queue. When the arbiter is programmed for sequence mode, frames from the EFPQ 36 which belong to the same command context as the previous frame 15 may be given priority. Thus, once a context begins, all frames belonging to it may be transmitted. In such a case, at the end of each context (or when the queue is empty), a frame from the EFIQ 34 may then be transmitted.

## 10. Egress Error Handling

Error handling may be accomplished by a combination of hardware error 20 detection and software error recovery procedures. The following will describe one embodiment of the hardware detection capabilities of the ASIC 10 egress path.

Each POS frame received by the ASIC 10 will typically contain a Frame CRC checksum. When an error is detected in this checksum, a status bit may be set in the segment of the EPB 40 that received the payload, according to one embodiment. The

manner in which the error may be handled is dependent (at least in part) on whether the frame header was routed to the Pass-Through Path or to the Internal Path.

If the header was routed to the Internal Path, the PRC 20 may be notified of the arrival of the frame after the payload has been fully received. In this

5 embodiment, the PRC 20 would check the receive status before processing the payload. If this check reveals that a receive error occurred, a software recovery procedure may be called. In one embodiment, part of the software recovery procedure would include returning the EPB 40 segment to the free pool, and releasing the HQM 24 entry to the hardware.

10 If the header was routed to the Pass-Through path, the PRC 20 may be notified of the arrival of the POS frame after the header is received, but while the payload is still in transit. Upon notification of the arrival of the POS header, the PRC 20 may create an FC header in an entry in the EFPQ 36 and release the entry to the hardware. This will normally occur before the POS CRC error is detected.

15 In order to handle this situation, the hardware that assembles the outgoing FC frames may be designed to examine the receive status field of the EPB 40 segment before it initiates the FC frame. If the status field indicates that a problem was encountered while receiving the POS frame, in one embodiment the state machine may transfer this status information to the entry in the EFPQ 36, turn the 20 entry over to the software, and halt without outputting the FC frame. The software may then decide how to recover from the error. In one embodiment, part of the recovery procedure would include returning the EPB 40 segment to the free pool and returning the FC HQM 24 entry to the hardware.

If Cut-Through mode is enabled, the system may start sending out FC frame 25 before the POS CRC error is detected. Such an error will typically be detected,

however, before the end of the FC frame has been transmitted. When this occurs, the hardware will end the FC frame with an EOFni (End of Frame, Normal Invalid), according to one embodiment. In another embodiment, the status field of the entry in the FC HQM 24 may be updated with information about the error, the entry 5 turned over to the software, and the hardware state machine halted. It should be appreciated that the software may then decide how to recover from the error.

Moreover, an additional hardware feature may be provided to help minimize the software recovery process. In one scenario, the frame with the CRC error advanced to the head of the EFPQ 36 before the software became aware of the error. 10 By that time, the HQM 24 could have contained headers of additional frames belonging to the same context. Furthermore, these frames could be interleaved with frames from other contexts. In order to allow the PRC 20 to easily purge frames belonging to a specific context from the HQM 24, a 'skip' bit may be provided in each entry in the HQM 24. When an error is detected, the PRC 20 can examine each 15 subsequent entry in a particular queue and set the skip bit in each frame it wants to purge. In one embodiment, this may be done before the PRC 20 re-enables the hardware. Once re-enabled, the hardware may process the HQM 24 in order, beginning with the entry after the one with the error. Thus, in this embodiment, each time an entry in which the skip bit set reaches the head of queue, its contents 20 may be ignore, the entry returned to the software, and the next entry processed.

Errors may also be encountered by the Egress Fibre Control (EFC) 44 module while sending FC Frames out on the wire. Such errors may be posted in the HQM 24 entry which originated the frame. After each FC frame is completed, either successfully or un-successfully, the HQM 24 entry that originated the frame may be 25 returned to the software. The PRC 20 may then examine the status field of the entry and if required, take appropriate recovery action.

One additional error condition may occur if Cut-Through mode is improperly set up. An error (e.g., ‘buffer under run’) can occur when a frame is being simultaneously received on the POS interface 12 and sent out on the Fibre 14. The error occurs if the speed on the sending side is greater than the speed on the

5 receiving side and the buffer runs out of data to send. If this occurs, the logic that generates the FC Frame may terminate the frame with an EOFni (End of Frame, Normal Invalid). The status field of the FC HQM 24 entry that originated the frame may then be filled in with information indicating the action taken, and the entry may be turned over to the software. In one embodiment, the processing of FC

10 frames from the Pass-through path is then halted. The software then has the option of re-transmitting the frame using the original HQM 24 entry, re-transmitting it using a new HQM 24 entry, or executing a recovery protocol.

### C. Frame Ingress

Each frame that is received from the Fibre Channel 14 may be routed to either

15 the “Pass-Through Path” or the “Internal Path.” In one embodiment, the Pass-Through Path is for frames containing payloads that will be sent out on the POS interface 12, while the Internal Path is for frames whose payload contains initialization, configuration or control information that will be used by an internal processor (e.g., PRC 20), but not sent out on the POS interface 12. In one

20 embodiment, the path to which the frame is routed is based on the contents of the R\_CTL field in the FC frame header.

After the routing decision has been made, the frame header may be stripped from the payload and stored in an entry in one of the two ingress FC Header Queues, according to the path (Pass-Through or Internal) that has been chosen. A

25 programmable number of bytes from the payload may also be stored along with the

header in the selected Header Queue entry. In the embodiment of Figure 1B, the two ingress FC Header Queues are the IFIQ 30 and the IFPQ 32.

In one embodiment, the header of the incoming FC frame is compared to the header of the most recent FC frame that was routed to the same path. If certain  
5 fields match, a bit in the status field of the FC HQM 24 entry may be set indicating that the frame belongs to the same context and is sequential.

The payload may then be separated from the frame and stored in the next available segment of the IPB 38, according to one embodiment. As will be described below in Section II, the IPB 38 may be comprised of a payload buffer and FIFO  
10 which may be used to provide out-of-sequence release of payload buffer segments and/or delivery of higher priority frames ahead of lower priority frames may be done using one or both of these data queues. Similarly, the IPB 38 module may also provide the capability of reordering out-of-order frames by, for example, holding an out-of-order frame until the expected in-order frame arrives.

15 A handle indicating which payload segment was used may also be stored in the FC HQM 24 entry that received the FC frame header. While in one embodiment the PRC 20 is notified that a frame has been received after the entire payload had been received, in another embodiment, this notification may occur before the entire payload has been received.

20 It should be appreciated that the PRC 20 may undertake a variety of operations at this point. The PRC 20 operation may be dependent upon several factors, including the path and contents of the frame, whether initialization has been completed, and in the case of an FCP frame, whether a command context already exists. Moreover, the PRC 20 may undertake a frame Pass-Through operation  
25 and/or an Internal Frame operation, as will now be described.

1. Pass-through Frame Operation

If the frame was routed to the Pass-Through path, the header would have been placed in an entry in the IFPQ 32, according to one embodiment. When the entry is turned over to the PRC 20, the PRC 20 may examine it and write the information necessary to create a suitable POS frame header in the next available entry in the IPPQ 28. A payload handle may also be copied from the FC HQM 24 entry to the POS HQM 24 entry. In another embodiment, if the frame belongs to the same context as the previous frame, it may use a mask field in the POS HQM 24 entry to tell the hardware to reuse portions of the previous POS frame header.

5 After the PRC 20 has finished setting up the outgoing POS frame header in the IPPQ 28, it may release the entry to the hardware. In one embodiment, this is done by setting a bit in the entry's control word. When the entry reaches the head of the queue, the hardware may automatically assemble the POS frame and send it out on the POS interface 12.

10 After the PRC 20 turns the entry in the IPPQ 28 over to the hardware, it no longer needs the entry in the IFPQ 32. Thus, in one embodiment the IFPQ 32 entry is released to the hardware for use by another Ingress FC frame by setting a bit in the entry.

15 When the entry in the IPPQ 28 reaches the head of a given queue, the hardware may then assemble a POS frame and send it out on the POS interface 12. When this has been completed, the completion status may be put into the outgoing HQM 24 entry that originated the frame, and the entry turned over to the software. Moreover, the payload buffer segment may be returned to the free pool, or it may be returned by the PRC 20 after the processor checks the completion status in the HQM 20 entry.

## 2. Internal Frame Operation

If the frame was routed to the Internal Path, the payload may be used by an internal processor (e.g., PRC 20). In one embodiment, a programmable number of payload bytes are available to the PRC 20 in the IFIQ 30, which may also be

5 accessible to the PRC 20 in zero-wait-state memory. In another embodiment, additional payload bytes may be examined by the PRC 20 via the F\_LIO bus 42.

After the PRC 20 has completed processing the FC HQM 24 entry, it may then return the entry to the hardware by setting a bit in the control word of the entry. The payload buffer segment may also be returned to the free pool by writing the

10 segment's handle to a register (e.g., "Payload Segment Release Register").

## 3. Special Payload Buffer

If the embedded processor (e.g., PRC 20) needs to generate an ingress POS frame, it may do so using the IPIQ 26 and the Special Payload Buffer. In one embodiment, the Special Payload Buffer is a single segment buffer consisting of a

15 predetermined number of bytes (e.g., 512 bytes) and resides in zero-wait-state processor memory. It should, however, be appreciated that other buffer configurations may also be used.

It should also be understood that the use of the Special Payload Buffer is optional, and will typically be used where the payload of the frame is too large to fit

20 into the spare bytes in the Header Queue entry. By way of a non-limiting example, when a nominal configuration of 128 bytes per Header Queue entry is used, there are 96 bytes available in each HQM 24 entry for a POS header and POS payload. If the total number of bytes of the frame to be sent is 92 or less, the entire frame can be put into an HQM 24 entry. Otherwise, the Special Payload Buffer may be used.

After the PRC 20 has put the required information into the HQM 24 entry and Special Payload Buffer, it may then turn the frame over to the hardware by setting a bit in the HQM 24 entry. In one embodiment, the hardware will queue the entry and send the frame out on the Fibre 14 when the entry reaches the head of the 5 queue.

#### 4. Optional Headers

When an FC frame is received, the FC header may be separated from the payload and stored in one of the two ingress FC Header Queues (Internal or Pass-Through). In one embodiment, a programmable number of additional bytes from 10 the FC frame are also stored in the Header Queue entry (e.g., HQM 24 entry). In another embodiment, the complete payload (everything after the FC header) may be stored in the next available segment of the IPB 38. If the bytes following the FC header contain an optional header, it may be located in the beginning of the payload buffer segment, as well as in the HQM 24 entry. In one embodiment, the PRC 20 15 may examine the optional header by reading it from the HQM 24 entry.

If the payload is to be forwarded to the POS interface 12, the PRC 20 may choose to exclude the optional FC header from the POS frame. In one embodiment, this is done by indicating the length of the optional header in a field (e.g., the “segment offset” field) of the ingress POS header queue entry that it generates for 20 the frame. When the payload is transferred, the hardware may then skip the number of bytes indicated by this field when it takes the payload from the IPB 38.

#### 5. Raw Frames

A frame that has been received on the Fibre Channel 14 may be fully encapsulated into a POS frame and sent out on the POS interface 12. In one

embodiment, there are two modes available to accomplish this operation, as will now be described.

The first mode, according to this embodiment, is a dedicated raw frame mode. When the Ingress Fibre Control (IFC) 46 logic is programmed for this mode,

5 each frame that is received from the Fibre 14 may be put into the IPB 38 in its entirety, including FCBB characters for the SOF and EOF characters of the frame. While it should be appreciated that less than the entire frame may be sent into the IPB 38, for illustrative purposes the following discussion assumes that the entire frame is put into the IPB 38.

10 In addition to being put into the IPB 38, the FC header may also be placed into an entry in one of the ingress FC header queues (e.g., IFIQ 30 and/or IFPQ 32). From this point on, the frame may be processed in the same manner as a normal Pass-Through frame. In one embodiment, the PRC 20 creates a POS header in the next available entry in the IPPQ 28, copies the payload segment handle to the queue

15 entry, and releases the entry to the hardware. When the entry reaches the head of the queue, the hardware may encapsulate the entire FC frame in a POS frame and send it out on the POS interface 14.

The second mode, according to this embodiment, is the interleave mode. In one embodiment, this mode allows raw frames to be interleaved with normal frames. In this mode, the hardware need not know in advance if an incoming FC frame will pass through as a raw frame, or if only the payload will be sent out on the POS interface. Moreover, in this mode the FC frame may be received in the same manner described above Section I.C.1.

After the PRC 20 has been notified of the arrival of the frame, it creates a POS header in the next available entry in the IPPQ 28 and copies the payload handle to

the entry, according to one embodiment. The PRC 20 may then determine if the frame should be treated as a raw frame or as a normal frame.

If the frame is to be treated as a raw frame, in one embodiment the following additional steps are performed before the POS HQM 24 entry is turned over to the  
5 hardware:

First, the PRC 20 copies the FC header from the FC HQM 24 entry to the POS HQM 24 entry. In this embodiment, the FC header may be written to the spare byte locations that immediately follow the POS header. The length of the FC header may then be written to a field (e.g., the `hdr_size` field) in the POS HQM 24 entry. This  
10 field can be used to tell the hardware that additional bytes (the FC header) will be taken from the POS HQM 24 entry after the POS header has been transferred, but before the payload is transferred.

Next, the PRC 20 copies the FC CRC checksum from the entry in the FC HQM 24 to the entry in the POS HQM 24 entry, according to one embodiment. In another  
15 embodiment, the PRC 20 may then tell the hardware to transfer this field after the payload by setting a bit in a field of the POS HQM 24 entry. In one embodiment, the bit that is set is the `imm-payld` bit in the `payld_src` field. In yet another embodiment, the PRC 20 may also indicate the length of the CRC checksum in the `imm_payld_size` field of the POS HQM 24 entry.

20 After completing these steps, the PRC 20 may then turn the entry in the POS HQM 24 entry over to the hardware. In one embodiment, when the entry reaches the head of the queue, the hardware builds the POS frame as follows: First, the POS header is generated using data from the POS HQM 24 entry. Second, the FC Header is transferred from the POS HQM 24 entry. Third, the FC payload is transferred

from the payload buffer segment. Fourth, the FC CRC is transferred from the POS HQM 24 entry. Finally, the generated POS frame CRC is transferred.

## 6. Cut-Through and Store-and-Forward Modes

In embodiment, ASIC 10 may provide two modes of operation. With the first 5 mode, referred to herein as the Store-and-Forward mode, frames are received in their entirety from the Fibre Channel 14 before they are sent out on the POS interface 12. Alternatively, the Cut-Through Mode described in previously-referenced co-pending U.S. Patent Application Serial Number 10/435,214 may be used. As 10 described therein, after a frame header and a programmable number of payload bytes have been received on the Fibre Channel 14 in this mode, the frame may be output on the POS interface 12. Thus, receiving and sending operations may overlap. In one embodiment, Cut-through mode may be enabled on a frame-by-frame basis.

## 7. Arbitration

15 In one embodiment, ingress POS Frames may originate in either the IPPQ 28 or the IPIQ 26. At any point in time, there may be multiple POS frame headers in each of these queues waiting to go out on the POS interface 12.

Within each queue, POS frames will be output in the order in which they 20 were released to the hardware, according to one embodiment. However, the same principle need not apply between queues. For example, frames that are waiting in one queue may be delayed while newer frames in the other queue go out on the POS interface 12.

In one embodiment, the arbitration algorithm has two settings: 'ping-pong' and 'sequence'. When the arbiter is programmed for ping-pong mode, ingress POS

frames may be taken from the IPPQ 28 and the IPIQ 26 in alternating order, one at a time from each queue. When the arbiter is programmed for sequence mode, frames from the IPPQ 28 which belong to the same command context may be given priority. Thus, once a context begins, all frames belonging to it may be transmitted in an uninterrupted fashion. In such a case, at the end of each context (or when the queue is empty), a frame from the IPIQ 26 may then be transmitted.

## 8. Ingress Error Handling

As with the Egress path, Ingress error handling for may be accomplished by a combination of hardware error detection and software error recovery procedures.

10 The following will describe one embodiment of the hardware detection capabilities of the ASIC 10 ingress path.

In one embodiment, each FC frame received by ASIC 10 will typically contain a frame CRC checksum and an EOF transmission word. When a checksum error or an EOFni (End of Frame, normal invalid) is detected, or any other Fibre-Channel-specific error is detected during the reception of a frame, a status bit may be set in the segment of the IPB 38 that received the payload. Moreover, the manner in which the error is handled may be dependent on whether the frame header is routed to the Pass-Through Path or the Internal Path.

If the frame is routed to the Internal Path, the PRC 20 may be notified of the arrival of the frame after the payload has been fully received. The PRC 20 may then check the receive status before processing the payload. In one embodiment, if the check reveals that an error condition occurred while receiving the FC frame, a software recovery procedure is called. It should be appreciated that the software recovery procedure called may include returning the payload buffer segment to the free pool, and releasing the HQM 24 entry to the hardware.

If the frame is routed to the Pass-Through Path, the PRC 20 may be notified of the arrival of the FC frame after the header is received, but while the payload is still in transit. In one embodiment, upon notification the PRC 20 creates a POS header in the IPPQ 28 and releases the entry to the hardware. While this will normally occur

5 before the POS CRC error is detected, it may also occur afterwards.

In order to handle this situation, the hardware that assembles the outgoing POS frames may be designed to also examine the status field of the indicated payload buffer segment before it initiates each POS frame. In such an embodiment, if the status field indicates that a problem was encountered while receiving the FC

10 frame, the state machine may transfer this status information to the POS HQM 24 entry, turn the entry over to the software, and halt without generating the POS frame. The software may then decide how to recover from the error. In one embodiment, the recovery procedure includes returning the payload buffer segment to the free pool and returning the POS HQM 24 entry to the hardware.

15 If, on the other hand, Cut-Through Mode is enabled, the hardware may start sending the POS frame out before the FC receive error has been detected. The error will typically be detected, however, before the end of the POS frame has been transmitted. When this situation occurs, the hardware may be given the option (programmable) of either corrupting the outgoing POS frame CRC, or indicating a

20 ‘Receive Frame’ error on the POS interface 12. In either case, the status field of the entry in the POS HQM 24 may be updated with information about the error. In one embodiment, the entry is also turned over to the software and the hardware state machine halted. In such a case, the software may then decide how to recover from the error.

In the example given above, the frame with the CRC error advanced to the head of the IPPQ 28 before the software became aware of the error. By that time, the queue could have contained headers for additional frames belonging to the same context. Furthermore, these frames could be interleaved with frames from other 5 contexts. In order to allow the PRC 20 to easily purge frames belonging to a specific context from the queue, a 'skip' bit may be provided in each queue entry. In this embodiment, when an error is detected the processor can examine each entry in the queue and set this bit in each frame it wants to purge. Thereafter, the queue may be processed in order, beginning with the entry after the one with the error. Thus, in 10 one embodiment, each time an entry with the skip bit set reaches the head of the queue, its contents may then be ignored, the entry returned to the software, and the next entry in the queue is processed.

## II. MANAGING PAYLOAD BUFFER SEGMENTS

### A. Overview

15 As mentioned above, one aspect of the invention is to process frames in an order that may be different than the order in which they were received. In one embodiment, the frames are processed by a protocol bridge where the incoming frames are Fibre Channel frames and the outgoing frames are sent out on a POS interface. It should be appreciated, however, that frames may be processed in 20 accordance with any networking protocol and that Fibre Channel and POS are but two such examples.

By way of providing a non-limiting example, a networking device (e.g., a protocol bridge, ASIC 10, etc.) may receive high priority FC frames from an FC interface. Such frames may contain control or link information. Moreover, these 25 frames may have been received after some lower priority frames, such as data

frames. It may be desirable to deliver the higher priority frames ahead of the lower priority frames over the POS interface. In one embodiment, this is done by providing out-of-sequence use of payload buffer segments. In another embodiment, delivering higher priority frames ahead of lower priority frames is done using a 5 single data queue.

Another aspect of the invention is for the networking device (e.g., protocol bridge, ASIC 10, etc.) to also have the capability of reordering out-of-order frames. In one embodiment, this requires the networking device to hold an out-of-order frame in its buffer until the expected in-order frame arrives. In such a case, the 10 frames may be released over the POS interface in a different order than the order in which they were received over the FC interface. In one embodiment, the reordering of out-of-order frames is accomplished using an on-chip processor.

## B. Free Pool FIFO

One implementation of the invention uses a Free Pool FIFO (“FIFO”) to keep 15 track of the order in which the payload buffer segments are available. In one embodiment, the FIFO provides a segment index for writing data into an associated payload buffer and also captures the segment index when the segment is being released back to the free pool by the reading block or by the embedded processor (e.g., PRC 20). In a further embodiment, the a single data queue (e.g., EPB 40 or IPB 20 38) is comprised of the FIFO and the payload buffer.

After power-on reset, the FIFO is cleared and no segments are available for reading or writing, according to one embodiment. The FIFO may then be initialized. Following FIFO initialization, the associated buffer (e.g., the payload buffer) may be available for use. In one embodiment, the first segment allocated is the first segment 25 that gets filled by incoming data. It should be appreciated that the write source may

provide an indication when the write transaction is completed. In one embodiment, this indication is a signal that causes a read pointer of the FIFO to increment to the subsequent FIFO entry. In another embodiment, the next entry in the FIFO (i.e., after the read pointer has been incremented) is displayed on a data-out bus. This

5 next entry may correspond to the next buffer segment to be filled by the next burst of incoming data.

When a segment of data is read out of the payload buffer, the reading block may provide a "release" indication (e.g., signal) when the read transaction is completed. In one embodiment, an indication from the reading block may cause a

10 write pointer of the FIFO to increment. In another embodiment, the released segment's index value may be written in the FIFO. While numerous processes for releasing data segments from the buffer may be used, in one embodiment the processor releases segments by (a) writing the segment index to a register, (b) setting the segment release request bit, and (c) polling an acknowledgement bit that

15 indicates the segment has been added to the free-pool.

Referring now to Figures 2A-2E, in which one example of the interaction between and data flow through a FIFO and an associated payload buffer data is depicted. While in the embodiment of Figures 2A-2E the payload buffer is comprised of sixteen segments, it should equally be appreciated that the buffer may

20 be comprised of more or fewer segments.

Figure 2A illustrates one embodiment of a payload buffer and corresponding FIFO after initialization. As shown in Figure 2A, the FIFO is populated with segment index values of SEG #0 in location 0 and sequentially increasing to SEG #15 for location 15. Both the write pointer (WR\_PTR) and the read pointer (RD\_PTR) are positioned at location 0 which contains the segment index value Seg #0, which will

25

be the first segment of the buffer that will get filled with incoming data, according to this embodiment. As each segment is filled, the write source may provide a “done” signal (or other indication), thereby incrementing the RD\_PTR of the FIFO to the next entry. Figure 2B depicts the payload buffer and FIFO after frames #0 - #3 have 5 been received. In particular, as the first four segments of the buffer were filled with frames #0 - #3, the RD\_PTR was correspondingly incremented and the index values were written to the current FIFO entry (e.g., the FIFO entry that was being pointed to by the RD\_PTR at the time the payload buffer segment was being filled).

Figure 2C depicts one embodiment of the payload buffer and FIFO after 10 frames #0 - #3 have been read from and transmitted out of the payload buffer in random order. In this example, frame #3 was transmitted first, at which point the released segment’s index value (i.e., SEG #3) was written to the FIFO at the then-current location of the WR\_PTR (i.e., location 0). Thereafter, the WR\_PTR was incremented. Frame #1 was thereafter transmitted, with the corresponding segment 15 index value being written to the FIFO and the WR\_PTR incremented again, followed by the same procedure for frames #0 and #2. At this point, the WR\_PTR and RD\_PTR are located at the same entry of the FIFO (as shown in Figure 2C), indicating that all received segments have been released (e.g., transmitted).

The example of Figure 2A-2C continues with Figure 2D, which depicts the 20 payload buffer and FIFO after frames #16 - #19 have been received. That is, Figure 2D assumes that frames #0-#15 were previously received and subsequently released. Specifically, in this example frame #16 was received first and placed in SEG #3, since that was the segment index value of the then-current location of the RD\_PTR (i.e., location 0). The RD\_PTR was then increment to the next FIFO entry (i.e., location 1). 25 In one embodiment, this incrementing is done after the write source provides an indication that the operation is complete. The next frame received (i.e., frame #17)

was then placed in SEG #1, again based on the segment index value of the current RD\_PTR position. This process continued until, as shown in Figure 2D, frames #16-#19 were all received.

At this point frames #16-#19 may be transmitted out of the buffer in any order. In the embodiment of Figure 2E, frame #17 was transmitted first. Upon receiving an indication that the transmission process was complete, WR\_PTR was incremented after writing that segment's index value (i.e., SEG #1) to the FIFO at the then-current location of the WR\_PTR (i.e., location 0). Thereafter, the next frame to be released was frame #19. Again, after writing that segment's index value (i.e., SEG #2) to the current location of the WR\_PTR, the WR\_PTR was incremented. This process continued until the last frame (i.e., frame #18) was transmitted out.

Figure 3 depicts one embodiment of the EPB 40 and IPB 38 modules consistent with the principles of the invention. Beginning with the EPB 40 module, As shown in the embodiment of Figure 3, the EPB 40 module contains an Egrs Payload Buffer 310 and an Egrs FIFO 320. In one embodiment, the Egrs Payload Buffer 310 and Egrs FIFO 320 may be used to manage data frames/payloads as described above with reference to Figures 2A-2E. In particular, the Egrs FIFO 320 may be initialized and populated with segment index values that correspond to segments of the Egrs Payload Buffer 310. Although not shown in Figure 3, a WR\_PTR and RD\_PTR may be associated with the Egrs FIFO 320 (as described above with reference to Figures 2A-2E) to keep track of the order in which payload buffer segments are filled and subsequently released (e.g., after segment data has been transferred). In addition, support logic 330 may also be included in the EPB 40 module to manage the interaction between the Egrs Payload Buffer 310 and Egrs FIFO 320, as well as the data flow to the EFC 44 module and from the Egrs POS Control 48 module.

Continuing to refer to Figure 3, IPB 38 module may also contain a payload buffer (i.e., Ingrs Payload Buffer 340) and a FIFO (i.e., Ingrs FIFO 350). As with the EPB 40, the Ingrs Payload Buffer 340 and Ingrs FIFO 350 of the IPB 38 module may be used to manage data frames/payloads as was described above with reference to Figures 2A-2E. In particular, the Ingrs FIFO 340 may be initialized and populated with segment index values that correspond to segments of the Ingrs Payload Buffer 350. Although not shown in Figure 3, a WR\_PTR and RD\_PTR (as described above with reference to Figures 2A-2E) may be associated with the Ingrs FIFO 350 to keep track of the order in which payload buffer segments are filled and subsequently released (e.g., after segment data has been transferred). Support logic 360 may also be included in the IPB 38 module to manage the interaction between the Ingrs Payload Buffer 340 and Ingrs FIFO 350, as well as the data flow from the IFC 46 module and to the Ingrs POS Control 50 module.

Figure 4 is one embodiment of a process 400 for managing payload load buffer segments in accordance with the principles of the invention. While in one embodiment, process 400 is carried out using ASIC 10, it should equally be appreciated that other networking devices may also perform the operations of process 400. It should be appreciated, however, that the following discussion of Figure 4 is intended to be a generalized discussion that is independent of whether the EPB 40, the IPB 38 or some other payload buffer module is involved. Moreover, references to incoming and outgoing data frames are not intended to correspond to the ingress and egress data paths discussed above in Section I. Rather, incoming data frames may be received by the networking device along either the ingress path or the egress path. Similarly, references below to outgoing data frames are simply those frames being transmitted out of the networking device, whether it be along the ingress path or the egress path.

Process 400 begins at block 405 with the initialization of the FIFO, which may be either the Egrs FIFO 320 or the Ingrs FIFO 350. In one embodiment, the process of initializing the FIFO involves clearing the FIFO such as, for example, by performing a power-on reset. Thereafter, the FIFO's entries may be populated with 5 the payload buffer's index values. At this point, the status of the payload buffer and FIFO would be as previously shown in Figure 2A.

Once the FIFO has been initialized, process 400 may proceed to block 410 where a determination is made as to whether there is an incoming frame. In one embodiment, the incoming frame is an FC frame that is received over Fibre Channel 10 14. Regardless of the type of frame being received, in the case of an incoming frame process 400 will continue to block 415 where the frame is received by the networking device (e.g., ASIC 10, protocol bridge, etc.). In one embodiment, the frame is an FC frame and is received by ASIC 10 as described above in Section I.

When the received frame includes a payload, process 400 may then continue 15 to block 420 where a storage location for the payload is selected. In one embodiment, the payload location is a segment of the payload buffer. In a further embodiment, the payload segment is selected based on the segment index value being pointed to by the FIFO's read pointer (e.g., RD\_PTR). Once the storage location has been selected, the received data frame may then be written to the 20 selected payload buffer location (block 425). In one embodiment, the write source may also provide a indication that the storing operation is complete. Thereafter, the read pointer (e.g., RD\_PTR) may be incremented to the next FIFO entry (block 430). Process 400 may then loop back to decision block 410 to determine if there is another incoming frame.

In addition to the aforementioned, after initialization process 400 may also proceed in a parallel processing path to decision block 435 to determine if there is an outgoing data frame. In one embodiment, this determination is based on detecting a signal (e.g., read enable signal) from the logic that is responsible for reading the 5 outgoing data frame from the buffer segment. In one embodiment of the ingress path, the IFC module 46 determines when to send an outgoing data frame. Once a determination is made that a frame should be sent, a read enable signal may be asserted. For the egress path, the EFC module 44 may make the decision when to send an outgoing frame. This decision may be based on the buffer-to-buffer credits 10 that it has received from a device connected to Fibre Channel 14, according to one embodiment. Again, a read enable signal may be asserted to indicate that a frame should be read from the buffer.

If a determination is made that there is an outgoing data frame, process 400 continues to block 440 where the requested data frame is transferred from its 15 payload buffer location. Thereafter, at block 445 process 400 write the segment index value for the transferred data frame in the FIFO entry currently being pointed to by a write pointer (e.g., WR\_PTR). In one embodiment, the segment index value is written to the designated FIFO entry only after receiving an indication that the data frame transmission process was completed. Once the segment index value for 20 the transferred data frame has been written to the corresponding FIFO entry, process 400 may then increment the write pointer to the next FIFO entry. Process 400 is then ready to loop back to block 435 to process the next outgoing data frame.

In this fashion, a Free Pool FIFO may be used to process frames in an order that may be different than the order in which they were received, such as processing 25 higher priority frames ahead of lower priority frames. Similarly, received frames may be reordered before being transmitted.

While certain exemplary embodiments have been described and shown in the accompanying drawings, it is to be understood that such embodiments are merely illustrative of and not restrictive on the broad invention, and that this invention not be limited to the specific constructions and arrangements shown and described,

5 since various other modifications may occur to those ordinarily skilled in the art.